



The Ultimate Guide to **Conversion Rate Optimization** in Ecommerce



Picture This.

You've chosen the best **machine learning (ML)-powered ecommerce search solution** for your ecommerce website. You've also conducted a thorough **audit** of your site search performance, finding that the products returned by your search engine are highly relevant, that type-ahead **suggestions** are very helpful, and that **facets** allow customers to seamlessly drill down into results.

To avoid showing only static category pages to the two-thirds of shoppers who use navigation for product discovery, you've enhanced those pages with an AI-optimized product sort order for each based on customer behavior and product performance.

Congratulations: you've now completed the first step towards full optimization of the conversion rate on your ecommerce website. There's also some slightly less good news, though, as there's still much more that you could — and should — be doing.

Don't worry. In this guide, we'll outline how to unlock further value from your investments by following three key principles.


1. **Leverage synergies between touchpoints and capabilities**
2. **Harness the full power of analytics**
3. **Test. Optimize. Repeat.**



Principle 1 Leverage Synergies Between Touchpoints and Capabilities

Ecommerce product search is critical to increasing your sales conversion rate. However, it can become an even more powerful component of your complex digital experience ecosystem when it works flawlessly with other mechanisms, functionalities, and points of interaction that facilitate product discovery, such as merchandising and recommendations.

We'll discuss some of these harmonies between search, recommendations, and also merchandising capabilities such as product badging, tailored offers, and content.



Ecommerce search and recommendations are important touchpoints


Ecommerce search is a crucial touchpoint driving a very large portion of retailers' online revenues. On average, **30%** of ecommerce visitors use on-site search and are known to convert between five to six times higher than non-site search visitors.

The increasing importance of product search is clearly visible when appreciating the evolution of Amazon's search box over the past 20 years. Check out Amazon.com from 2001. The search box was only big enough for a maximum of two words. Why? Amazon didn't trust their search box and neither did the customers. But search has come a long way. Today, Amazon's search box is front and center. Customers trust it and, when they use it, they (usually) get relevant results.

On the flip side, customers have come to love product recommendations as well. For example, consider **McKinsey's claims** that 35% of what consumers purchase on Amazon and 75% of what they watch on Netflix come from product recommendations. While **some have suggested** that these estimates should be more conservative, they nevertheless give an idea of the importance of product recommendations.

It's not about whether search is better than recommendations, or vice versa. Rather, you should focus on making the most out of how these tools complement each other to optimize conversion rates.


Amazon's search box evolution over the past 20 years



Tools to optimize conversion rates

Zero-result recommendations

Zero-results pages are a big problem for conversion rates. When customers try a search and get no results, there's a high risk of site abandonment. Most no-result pages typically offer few or no paths for customers to get back on track, making it a UX dead end.




Whatever the reason for a customer reaching a zero-results page, when there's little or no helpful guidance or content on the page, customers have two options: either think up a whole new product-finding strategy or abandon the site (typically to Google or Amazon).

Zero-result recommendations help turn this potential disaster into an opportunity for product discovery. Though users may not find the exact products they were searching for, zero-result recommendations are an effective way of pointing them towards other products they may be interested in, an opportunity often missed by most ecommerce websites today. Read [this](#) for more details on how to leverage zero-result product recommendations.

Session-aware recommendations

Imagine I've searched for some golf pants on an ecommerce website, which I've promptly added to my cart. The ecommerce site I'm on starts providing me with golf related items automatically. This is an example of **so-called** session-aware recommendations, which is a game-changer that helps ecommerce players deliver unique, 1:1 personalised shopping experiences without the need for high volumes of data or logged users.


Session-aware recommendations work by understanding user intent based on interactions observed in an ongoing session and analyzing the search and browsing behaviour in real-time.



These recommendations need to remain relevant throughout the user's session by also detecting intent changes and adjusting accordingly along their journey. You wouldn't want to keep seeing golf accessories once you start shopping for your next tennis outfit!

Tailored offers and badges


When it comes to optimizing your ecommerce conversion rate, don't underestimate the importance of adding relevant, tailored offers and badges that focus on social proof, scarcity, and urgency.



For example, you can both increase loyalty and repeat purchase by predicting when customers will need to stock up on their favorite products and then auto-serve personalized replenishment prompts. A returning customer would then see a 'it's time to stock up' notification, like was done by our customer **MAC**, who witnessed how 67% of those banners would be clicked, thus promoting further purchases on their website.

Another opportunity to increase conversion lies in understanding shopper potential 'exit' behavior. For example, a customer highlighting the product name with their mouse cursor is likely going to copy-paste it in another tab in order to price-shop.

Coveo's technology allows you to not only detect and understand exit behavior via shoppers' signals, but also to serve a relevant offer that compels them to stay and purchase.



Customer highlighting the product name in order to price-shop

See how Express Gifts took advantage of our technology.

Plants
Bulbs
Indoor plants
Pots
Inspiration
Outdoor living
Tool shed
Wildlife
Florist
Christmas
Offers

Home > Plants > Shrubs > Magnolia > *Magnolia* 'Susan'

VIDEOS

Viewed 14 times in the last 24 hours

Magnolia 'Susan'
magnolia

(1 review)
[Write review](#)

20% OFF mature shrubs

12 lt pot (0.8-1m)
£99.99 £79.99
 shipped within 2 weeks

Quantity
1

ADD

add to wishlist
 pot size guide →


Principle 2 Harness the Full Power of Analytics

There are a few key requirements to look for in any best-in-class analytics solution, from providing insights based on the metrics that matter most in an ecommerce setting to handling the complexity of today's customer journeys. Let's take a closer look.

Look at the big picture to gauge overall impact

Analytics need to provide details on a number of relevant ecommerce-specific KPIs. While you want to boost your conversion rate, a holistic approach is important. Measuring KPIs in isolation can be misleading and even have the opposite impact of what you were aiming for.

When tracking metrics for ecommerce sites, it's important to understand that metrics simply show symptoms. Different symptoms become visible through different metrics. You may see your Average Order Value (AOV) increasing, only to find that the Conversion Rate (CR) is tanking. Or alternatively, during some campaigns, you may have increases in CR due to some promotions and discounts, which are accompanied by a decrease in AOV.



That's why several companies look at **Revenue per Visit (RPV)** with great interest, as it merges both CR and AOV into one measurable metric without leaving any blind spots (i.e., $RPV = AOV \times CR$).

Also note that looking at Click-Through Rates (CTR) alone won't tell you the entire story about the success of your CR optimization initiatives, CTR is still a very important metric to track. For organizations where customers explore catalogs online but complete purchases in a physical store, it can be valuable to keep an eye not only on CR but also on CTR.

Don't underestimate the value of qualitative analyses


A dedicated site search analytics dashboard lets you conduct valuable qualitative analyses of visitor search behavior on your website based on their search journeys. Detailed walkthroughs on how to conduct these analyses can be found in Louis Rosenfeld's book *Search Analytics for Your Site*, but let's consider some examples here.

For instance, by looking at the list of queries used by your shoppers, you can get valuable ideas to fuel your merchandising and marketing initiatives. Noticing the popularity of thematic words like "healthy," "cheap," or "revolutionary" can drive your future campaigns. You can then use features such as **Coveo's advanced merchandising capabilities** to easily create clickable banners with those terms that get shoppers' attention and showcase special deals, thus improving your ecommerce CR.

You should also analyze your zero-result queries, as these can shine a light on issues and improvements for your ecommerce site searchers. Not offering the products that customers are looking for is one thing, but if queries are leading to zero results even though you do sell the products, then you may have a relevance problem.

The rise of omnichannel commerce and the end of shopping's boundaries means that looking at CR alone can occasionally lead to partly inaccurate conclusions.


Functionality such as typo-tolerant, as-you-type search suggestions that feed on machine learning are de-facto standards in today's market. Understanding where your search solution may be falling short can also serve as a basis to justify your case for investing in **AI-powered search capabilities**.



Ensure your analytics capture the entire customer journey

As we saw, customers engage with your ecommerce website across different touchpoints. There's rarely a conversion that occurs based on a single interaction. In fact, conversions typically stem from a mix of interactions with search, category listings, and product recommendations.

Path to conversion



Today's customer journeys are increasingly fragmented. Your analytics solution needs to handle their complexity, providing insights on behavior across multiple touchpoints. Yet, assessing performance in such a complex digital environment requires sophisticated ways of handling event attribution.

When determining which touchpoint plays the most important role in driving conversions, single-touch attribution models are not the best choice. Best-in-class solutions have recognized the need to move beyond simplistic last-touch and first-touch attribution models that assign 100% of the credit for a conversion to a single touchpoint in the customer journey.



Despite the difference between these two models (the most obvious being which touchpoint receives credit), both first-and last-touch attribution models fail to acknowledge the dramatic shift in shoppers' consumer journeys, implying that shoppers' decisions aren't influenced by a variety of interactions.

Your analytics solution therefore needs to embrace a "multi-touch" model — one that analyzes and acknowledges customer behavior throughout the journey to purchase.




Let's break this down

Multi-touch attribution refers to determining the value of each customer journey touchpoint that results in a conversion. In particular, we may want to determine the value of different on-site customer touchpoints (such as search, recommendations, and category listings) that contribute to a conversion event.

Imagine that a shopper has added a product to their cart after a search. Later, they added a second product to their cart, this time influenced by product recommendations, before converting. A multi-touch approach can easily split attribution equally to search and product recommendations, whereas a last-touch attribution model would only credit the last touchpoint of the customer journey for the end conversion.

Ultimately this would result in the impact of search on conversions being under-reported and the impact of recommendations on conversions being over-reported.

"Multi-touch" path to conversion




Principle 3 Test. Optimize. Repeat.

Securing best-in-class analytics won't release you from the need to embrace a culture of experimentation. You also need to constantly run experiments to test and optimize the experiences that your increasingly demanding customers receive.

Picture this: you just launched a new feature on your ecommerce search engine. You bet that conversion rates for ecommerce search will improve because of the feature's impact on your shoppers' experience. After a week, you look for some answers regarding the impact of that new and supposedly awesome feature.

The first and most intuitive place to look is the raw correlation. So you look at analytics and find that the conversion rate is higher than the one recorded prior to the new feature's release. It's also better than the conversion rate recorded during the same period the previous year.



Can you conclude that your new feature caused an increase in conversion rate? Unfortunately, raw correlations alone are rarely actionable. Correlation is not causation, and the complicating factor here is a set of other elements that may compete with the new feature in explaining the outcome. Economists call these confounding variables.



Suppose we observe that the sales of ice cream might be highest when the rate of drownings in city swimming pools is highest. Needless to say, this does not mean that ice cream sales are the cause for the increase in drownings, or vice versa. A plausible alternative explanation would be that a heat wave caused both spikes in ice cream sales and drownings in swimming pools. If this is the case, then the heat wave acts as a confounding variable — a cause common to both observations. In such a situation, a direct causal link cannot be inferred.

Although there are a number of econometric methods to get around confounding variables and credibly estimate causal relationships, A/B testing is the most robust way to prove causation.

Digital giants such as Google, Facebook, and Amazon have gotten to where they are by listening to their customers through rigorous data collection and [A/B testing](#). But all ecommerce companies should regularly conduct A/B tests and show different versions of an online experience to different sets of users. This is critical to understanding how to drive more conversions, as well as generally determine how and what to tweak to create a given outcome.

You can apply A/B testing to a broad range of site elements, such as landing pages, banners, or online advertisements. In the context of ecommerce search, different search experiences typically serve as the variants in an A/B test. By setting up a randomized percentage allocation of visitors to each experience, you can compare the impact of each. The best results are obtained by restricting the number of variations within an experience, using adequate sample sizes and concentrating on observed trends versus day-to-day fluctuations. For example, you can test deploying search personalization or specific machine learning models.

It's hard to overstate the importance of A/B testing, as this can be a powerful tool to prove causality. By randomly assigning an experience to some users and not others, we can precisely estimate the causal impact of the experience on the outcome we care about. In other words, by comparing two variants, such as a personalized product listing page and a non-personalized one, one variant can be declared to be more effective than the other as it leads to better conversion rates.

While the increasing adoption of A/B testing capabilities makes it easier to embrace this culture of experimentation, it is critical to keep in mind that A/B testing can prove causality only when this follows well-established, rigorous, scientific practices. Scientists use rigorous testing protocols to ensure they have a clear line of sight into what they are trying to prove or disprove. Digital tests need not be different.



The outside world often has a much larger effect on metrics than product changes do. Users can behave very differently depending on the day of week, the time of year, the weather (especially in the case of a travel company like Airbnb), or whether they learned about the website through an online ad or found the site organically. Controlled experiments isolate the impact of the product change while controlling for the aforementioned external factors.

Airbnb's data science team

Five recommendations to keep in mind when A/B testing

1

Start with a clear question and formulate a hypothesis

Of course, that implies selecting and monitoring meaningful Overall Evaluation Criteria (OEC), or quantitative measures of the experiment's objective. For example, your OEC might be the CR or the AOV. In proper test design, you should decide on the OEC you're going to look at before you execute an experiment. The more you're measuring, the more likely that you're going to see random fluctuations.

2

Have a big enough sample size

Often, A/B tests are conducted on too small of a sample. This means that there is not enough data to confidently say that one variation is more successful than another because not enough visitor traffic was tested. The best way to avoid this issue is to use a [sample size](#) calculator instead of guesstimates.

3

Ensure random testing

Making sure that a test is truly random is difficult. For example, as suggested before by the quote from Airbnb, one should always pay attention to day-of-week effects (you may have a different population of users on weekends than weekdays, so it's important that your experiment captures the normal business cycle). This can otherwise skew the results of an A/B test.



4

Focus on statistical significance

Merely comparing increases, drops, and general differences between KPIs isn't enough. You should always accompany your analyses with a measure of statistical significance. This is because when we conduct an A/B test, we are observing only one sample. If we took another sample or did another experiment, then the result would almost certainly vary, meaning that there is always some uncertainty in our results. The purpose of statistical inference is to estimate this uncertainty, and there are different frameworks available for this.

While Bayesian inference refers to statistical inference where uncertainty in inferences is quantified using probability, statistical practice during the 20th century and still early 21st has been dominated by a frequentist approach, which never uses or gives the probability of a hypothesis. (Here's a [quick rundown](#) of the two approaches' differences.) Theoretical disputes aside, under the dominant Null Hypothesis Significance Testing (NHST) paradigm, experimenters typically look at the [p-value](#) and follow the scientific standard to use a p-value less than 0.05 as indicator of statistical significance.

5

Keep A/B tests manageable

While it might be tempting to test many different variants at the same time, proceed with caution. A/B and multivariate tests are often used synonymously, but they are not the same thing. A/B testing simply compares A versus B versions of a customer experience using two different sets of users. Multivariate testing involves testing multiple versions of multiple elements at the same time to determine the winning combination. When you test many variations at the same time, however, you run into what is known as "cumulative alpha error." Basically, the more test variants you run, the higher the probability of false positives.

So Where Should You Start?

1 Choose which scenarios you would like to measure

Maybe you're wondering if changing the recommendation strategy on your ecommerce website's homepage from "Most Popular Bought" to "Most Popular Viewed" will convert more. Instead of simply switching the strategies on your recommendation carousel and running a pre-post analysis, you instead run an A/B test to measure the efficacy of this change.

Before enabling this test, make sure that you have everything in place to make an informed decision about which strategy you should keep.

2 Set up tracking for the relevant metrics

To ensure you're capturing the impact of your scenarios, you'll need to make sure you're tracking the right metrics. In this case, make sure you're tracking sessions and purchases to be able to calculate your conversion rate. These metrics can vary depending on which hypothesis you're trying to validate, so you'll need to make sure that you already have the required data tracking in place for the hypothesis you want to confirm.

Note that tracking the required events is not enough. To be able to report on the right attribution, you also need to make sure that all your events are rigorously identified with their origin to reflect on the multi-touch nature of customer interactions. This way, you'll be able to specifically measure the impact of your recommendation component instead of the overall experience. This reduces the chance for error in your test and increases the precision of your conclusions.

Finally, for this step, it's also important to make sure that you have a way to differentiate between the different versions of your A/B test. Either by identifying your visitors, sessions, or events with a specific value that can be used to filter on your reports.

This way, you'll be able to measure the distribution between your scenarios and ensure that both your test and reporting are yielding accurate results.

3 Complement business relevance with statistical significance

Measuring two versions of your site for any period of time will not necessarily guarantee that your test results will be significant enough to prove your hypothesis. As was stated earlier, it's important to define the required sample size based on the precision you want to achieve.

To do so, you can use a [sample size](#) tool to estimate how much data is needed to complete your test. You can then estimate the duration of the test by comparing the required volume of data with your normal site traffic.

You can always run shorter tests, but note that small samples mean low statistical power, and increase the likelihood that a statistically significant finding represents a false positive result.

4 Make informed decisions based on the outcome of your test

Assuming that everything went according to plan and you have run your A/B test for the right period of time to get the right level of statistical significance, you should now have results that can be interpreted to make an informed decision. However, just because you achieved better or worse results does not mean that the work stops here. There are always opportunities to improve and iterate quickly by constantly testing new approaches to improve on the overall experience.

A/B tests provide a quick and efficient way to validate your hypothesis and gain more knowledge around your customers' preferences.

Experimentation should be at the heart of every business, driving decisions that provide the best customer experience possible. In the end, this will improve both your business revenue and your customers' satisfaction.



Learn more about Coveo

Coveo is the world's leading cloud-based relevance platform. The Coveo Relevance Cloud™ uses applied AI to deliver relevant experiences in all digital interactions, from search to recommendations to personalization.

Contact us

